

A Unified Relevance Retrieval Model by Eliteness Hypothesis

Jagadeesh Gorla¹, Stephen Robertson² and Jun Wang¹
{j.gorla,j.wang}@cs.ucl.ac.uk, stephenerobertson@hotmail.co.uk

¹ University College London

² Microsoft Research, Cambridge

Abstract. In this paper, an Eliteness Hypothesis for information retrieval is proposed, where we define two generative processes to create information items and queries. By assuming two types of deterministic relationships between the eliteness of terms and relevance, we obtain a new theoretical retrieval framework. The resulting ranking function is a unified one as it is capable of making use of available relevance information on both the document and the query, which is otherwise unachievable by existing retrieval models. Our preliminary experiment on a simple ranking function has demonstrated the potential of the approach.

1 Introduction

A central problem in information retrieval modeling is concerned with the estimation of the probability of relevance between documents and queries. In existing probabilistic retrieval models, the probability of relevance is calculated under two views. The first view is conditioned by the given document. That is to correlate a given document with the information need properties of the users who would judge it relevant. The second view is conditioned by the given information need (query). That is to correlate the given information need with the information properties of those documents that they would be judged relevant [7]. The former is called the *document-oriented view*, including Maron and Kuhn's Probability Indexing and the language models [4, 5]; whereas the latter is called the *query-oriented view*, represented by the Robertson-Spärck Jones model [8]. The two views rely on fixing one variable and optimising the other, i.e. conditioned by the information need and tuning the document or the other way around, but not both [6]. Robertson et al [7] first identified this problem and proposed an initial theoretical framework of unified model to unify both the views. One of the advantages of such modelling is that it is capable of using available relevance information on both the document and the query.

In this paper, we formulate a novel unified retrieval framework based on the "Eliteness Hypothesis" proposed in [2]. The framework is validated by implementing and testing a simple resulting ranking formula using the Expectation-Maximization parameter estimation method [1].

The basic principle of a unified model is the following [6]: firstly describe the documents and queries using a set of properties, including the vocabulary

terms used to express them, and then the retrieval is performed by computing the relevance between document-query pairs as a function of description their properties. It should be possible to modify the descriptions of both the document and query properties once the relevance information has been available. Based on this idea, we define an Eliteness Hypothesis under the assumption that any given information items (document or query) can be described using a set of elite properties, and the relevance between a document-query pair is then dependent only on the elite properties and it can be thus estimated as a function of their elite properties.

In the following sections, we first discuss the concept of eliteness (Section 2) then we state the “Eliteness Hypothesis” for IR and define two deterministic relevance relations between a document-query pair, namely “Strict Identity” and “Logical inclusion”, under the hypothesis (Section 3). In section 4, we derive a unified probabilistic relevance ranking function for IR based on Eliteness Hypothesis by probabilistically modelinBased on thisg the eliteness of elite propeRelevance plays a critical role in information retrieval modelling. In many mathematical formulations, the probability of relevance is calculated either by correlating the documents with the information need properties of the users who would judge them relevant, or by correlating the users information needs with the information properties of those documents that they would be judged relevant. In other words, existing probabilistic retrieval models are capable of using the relevance information present either on query or document but not both. In this paper, we propose a theoretical framework to develop a unified probabilistic relevance model. The intention is to make use of the available relevance information on both the query and the document in order to compute the probability of relevance between them. Our development is based on the “Eliteness hypothesis”, which assumes that any document or a query can be described using a set of elite properties. The relevance between the document and query is dependent only their elite properties and these properties can be modified in the presence of relevance feedback. We derive probabilistic relevance ranking functions based on our hypothesis and as a preliminary study, evaluate one of simplified versions on the TREC-8 ad hoc task collection.rties. Furthermore, we derive two probabilistic ranking functions from unified relevance ranking function using “Strict identity” and “Logical inclusion” relationships along with some valid simplifying assumptions. For implementing a basic ranking function, in order to validate the proposed theoretical framework, we present an estimation methode for estimating the eliteness of elite properties by assuming the document collection is a 2-Poisson mixture [11] of elite and non-elite documents for a given elite property (Section 5). Finally, we present experimental results on the basic ranking function and show that it performs better than other influential ranking functions on TREC-8 ad hoc collection. We will also present some discussion about the basic ranking function relationship with the Inverse Document Frequency [?].

2 Eliteness

The idea of “eliteness” is to describe a set of elite properties (or concepts) represented by an information. In general, in the context of IR, we can classify elite properties into two different sets of properties:

1. properties that describe both information item (document) and information need (query). A typical example of these properties are the properties associated with the terms in vocabulary. These properties could be the natural language phrases used to express the information but in this paper we use only the elite properties that are associated with each term in the vocabulary. We name the properties associated with the single term properties as term-description properties.
2. properties that describe only document (e.g. page rank, url depth, etc) or only query.

We define the eliteness of an elite property for the given information, on the basis of the definition of eliteness introduced in [11, 2] as:

- the eliteness value of an elite property for the given information is binary, i.e. it is either elite (1) or non-elite (0). We can also say that the eliteness of a piece of information to a given elite property is a binary one.
- If the piece of information is *about* an elite property then its eliteness is elite for the property otherwise non-elite.

From the definition of eliteness, for any given elite property, some documents in the collection are “elite” and others are “non-elite”. Inspired by a 2-Poisson model proposed in [2, 11], we state our Eliteness Hypothesis as:

“Any information, expressed by an author or a user, can be described using a binary set of elite properties. If we know the eliteness value of each elite property for a given information item (document) and need (query), we can deterministically determine whether the item-need pair is relevant or not.”

We further assume that an author (or a user) will carry out the following process to express their information:

1. First, a user or author will choose a set of elite properties such that these properties can describe every aspect of the information that they want to express. The chosen elite properties are “elite” and the rest are “non-elite”.
2. Once the elite properties are chosen, an observable information item or need, is generated by a stochastic function of chosen elite properties. The uncertainty about the eliteness is injected during the generation process.

Our hypothesis suggests that the relevance between document d and query q is known exactly if their elite properties, represented by \mathbf{E} , \mathbf{F} respectively, are observed. Two deterministic relations are defined as follows:

A document d and query q are relevant if

- Strict identity: $\mathbf{E} = \mathbf{F}$, i.e. the eliteness of all the elite properties of d must be the same as the eliteness of the elite properties of q , or
- Logical inclusion: $\mathbf{F} \subset \mathbf{E}$, i.e. the elite properties of the document with eliteness “elite” must contain all the elite properties of the query whose eliteness is “elite”.

The other possibilities of defining the relevance under the hypothesis are stochastic rather than deterministic but would introduce an additional stochastic layer into the model. For simplicity, this paper only considers the deterministic relation and leave the stochastic one for future work. However, it should be noticed that the relevance between d and q will still be a probabilistic function given the generative process from Step 2 and the fact that the eliteness is not observable. The unification is achieved by modifying the eliteness probabilities of the document-query pair when the relevance information is available on the document and/or the query, which is otherwise not provided by existing retrieval models.

3 The Probabilistic Relevance Ranking Function

We are ready to derive a probabilistic ranking function. Let D be a random variable whose value is any possible document generated by an author. Similarly, Let Q be a random variable whose possible values are any possible query generated by a user. We use the lower case d and q to denote their particular instantiations respectively. Let $\mathbf{E} \in \{0, 1\}^k$ be a random binary vector over the space defined by the document’s elite properties, where k is the number of elite properties, e.g. one elite property for each word in vocabulary. Similarly, $\mathbf{F} \in \{0, 1\}^k$ is a random binary vector over the space defined by the query’s elite properties. Their elements are denoted as E_i and F_i respectively, where $i \in \{1, k\}$. $E_i = 1$ means “elite”, whereas $E_i = 0$ means “non-elite”.

Our eliteness hypothesis gives the probability of relevance $P(R = 1|d, q)$ as

$$P(R = 1|d, q) = \sum_{\mathbf{E}} \sum_{\mathbf{F}} P(R = 1, \mathbf{E}, \mathbf{F}|d, q) \quad (1)$$

where R is a binary random variable. $R = 1$ means relevant while $R = 0$ otherwise. Applying Bayes’ rule gives

$$P(R = 1|d, q) = \sum_{\mathbf{E}} \sum_{\mathbf{F}} \underbrace{P(R = 1|\mathbf{E}, \mathbf{F}, d, q)}_{\text{part one}} \underbrace{P(\mathbf{E}, \mathbf{F}|d, q)}_{\text{part 2}} \quad (2)$$

From the hypothesis, \mathbf{E} is independent of q and \mathbf{F} , \mathbf{F} is independent of d and \mathbf{E} and \mathbf{E}, \mathbf{F} are sufficient to determine the relevance between d, q . Therefore, by ignoring d and q and applying independence assumptions to Eq. (2), we get

$$P(R = 1|d, q) = \sum_{\mathbf{E}} \sum_{\mathbf{F}} \underbrace{P(R = 1|\mathbf{E}, \mathbf{F})}_{\text{part 1}} \underbrace{P(\mathbf{E}|d)P(\mathbf{F}|q)}_{\text{part 2}} \quad (3)$$

By applying Bayes' rule to Eq. (3) and removing constant $P(R = 1)$ we get

$$P(R = 1|d, q) \propto_R \sum_{\mathbf{E}} \sum_{\mathbf{F}} \underbrace{\frac{P(\mathbf{E}, \mathbf{F}|R = 1)}{P(\mathbf{E})P(\mathbf{F})}}_{\text{part 1}} \underbrace{P(\mathbf{E}|d)P(\mathbf{F}|q)}_{\text{part 2}} \quad (4)$$

where \propto_R denotes the rank equivalent. Here, we make an assumption, similar to term-independence assumption in [8], that the eliteness of an elite property is independent of other elite properties. We thus have

$$P(R = 1|d, q) \propto_R \sum_{\mathbf{E}} \sum_{\mathbf{F}} \prod_i^k \left(\frac{P(E_i, F_i|R = 1)}{P(E_i)P(F_i)} P(E_i|d)P(F_i|q) \right) \quad (5)$$

The Eq. (5) is a unified probabilistic relevance ranking function. It uses the information about each elite property between the document and query in the relevant set ($P(E_i, F_i|R = 1)$). Thus, the information about other relevant document-query pairs that share the eliteness of elite property is included, which is an essential component of a unified model [7].

In order to test the theory, we make the following assumption called *Query Eliteness Assumption* and derive relevance ranking functions from unified ranking function using the Strict Identity and Logical Inclusion relations.

Query Eliteness Assumption: We assume that the eliteness of each elite property is known for a given query, i.e. we know the binary vector \mathbf{F} . This assumption is similar to an implicit assumption made the query terms are elite to the query and other terms are non-elite in [11]. Let f_{q_i} denotes the eliteness value of F_i for each of term, namely $F_i = f_{q_i}$ where $f_{q_i} = 1$ if F_i is elite for q otherwise $f_{q_i} = 0$. Based on the our assumption, Eq. (5) can be written as

$$P(R = 1|d, q) \propto_R \sum_{\mathbf{E}} \prod_{i=1}^k \underbrace{\left(\frac{P(E_i, F_i = f_{q_i}|R = 1)}{P(E_i)P(F_i = f_{q_i})} P(E_i|d)P(F_i = f_{q_i}|q) \right)}_{\text{part 1}} \quad (6)$$

By applying Bayes' rule to part 1 of the Eq. (6), we get

$$\begin{aligned} & P(R = 1|d, q) \\ & \propto_R \sum_{\mathbf{E}} \prod_{i=1}^k \left(\frac{P(E_i|F_i = f_{q_i}, R = 1)P(F_i = f_{q_i}|R = 1)}{P(E_i)P(F_i = f_{q_i})} P(E_i|d)P(F_i = f_{q_i}|q) \right) \end{aligned} \quad (7)$$

Further factorizing Eq. (7) gives

$$\begin{aligned} & P(R = 1|d, q) \\ & \propto_R \sum_{\mathbf{E}} \prod_{\forall i: F_i=0} \left(\frac{P(E_i|F_i = 0, R = 1)P(F_i = 0|R = 1)}{P(E_i)P(F_i = 0)} P(E_i|d)P(F_i = 0|q) \right) \\ & \quad \prod_{\forall i: F_i=1} \left(\frac{P(E_i|F_i = 1, R = 1)P(F_i = 1|R = 1)}{P(E_i)P(F_i = 1)} P(E_i|d)P(F_i = 1|q) \right) \end{aligned} \quad (8)$$

From the Query Eliteness Assumption, we know that the value of each element in \mathbf{F} . That is there is no uncertainty in the eliteness of elite properties for the given q . So we have $P(F_i = 1|q) = 1$ if term i is in query q and $P(F_i = 0|q) = 1$ otherwise. By substituting the above values, Eq. (8) becomes

$$P(R = 1|d, q) \propto_R \sum_{\mathbf{E}} \prod_{\forall i: F_i=0} \left(\frac{P(E_i|F_i=0, R=1)P(F_i=0|R=1)}{P(E_i)P(F_i=0)} P(E_i|d) \right) \prod_{\forall i: F_i=1} \left(\frac{P(E_i|F_i=1, R=1)P(F_i=1|R=1)}{P(E_i)P(F_i=1)} P(E_i|d) \right) \quad (9)$$

From here onwards, we use Eq. (9) to derive a relevance ranking functions equivalent of the relevance based on Strict identity and Logical inclusion.

3.1 Ranking function under Strict identity

As defined under the Strict Identity relation, a document and a query is relevant if $\mathbf{F} = \mathbf{E}$. From the definition, we can assume that if we know that the eliteness of an elite property is “elite” for the given query (i.e. $F_i = 1$), then the probability that eliteness of same property is “elite” ($E_i = 1$) in the relevant set of documents, i.e. $P(E_i = 1|R = 1, F_i = 1) = 1$, as they have the same eliteness value. Equally, it is the same for the non-elite case where $P(E_i = 0|R = 1, F_i = 0) = 1$. Note that from this assumption, the score of any vector in \mathbf{E} of Eq. (9) is zero if $E_i = 0$ and $F_i = 1$ (or) $E_i = 1$ and $F_i = 0$ for at least one i . By substituting the above values in Eq. (9), we get

$$P(R = 1|d, q) \propto_R \prod_{\forall i=1: F_i=0} \left(\frac{P(F_i=0|R=1)}{P(E_i=0)P(F_i=0)} P(E_i=0|d) \right) \prod_{\forall i: F_i=1} \left(\frac{P(F_i=1|R=1)}{P(E_i=1)P(F_i=1)} P(E_i=1|d) \right) \quad (10)$$

where $P(F_i = 0|R = 1)$, $P(F_i = 1)$, $P(F_i = 0)$ and $P(F_i = 1|R = 1)$ in Eq. (10) can be removed as these terms do not effect the ranking order. We thus get

$$P(R = 1|d, q) \propto_R \prod_{\forall i=1: F_i=0} \left(\frac{P(E_i=0|d)}{P(E_i=0)} \right) \prod_{\forall i: F_i=1} \left(\frac{P(E_i=1|d)}{P(E_i=1)} \right) \quad (11)$$

Eq. (11) is a ranking function under the Strict identity relation with the Query Eliteness Assumption.

3.2 Ranking function under Logical inclusion

As defined under the Logical Inclusion relation, a document is relevant to a query if and only if $\mathbf{F} \subset \mathbf{E}$. We can thus assume $P(E_i = 1|R = 1, F_i = 1) = 1$

as all the elite properties with eliteness value “elite” for a query must have the same eliteness value in a relevant set of documents. That is, if a property $F_i = 1$ then $E_i = 1$ in relevant documents. By applying the above assumption, Eq. (9) becomes

$$P(R = 1|d, q) \propto_R \underbrace{\sum_{\mathbf{E}'} \prod_{\forall i: F_i=0} \left(\frac{P(E_i|F_i=0, R=1)P(F_i=0|R=1)}{P(E_i)P(F_i=0)} P(E_i|d) \right)}_{\text{part 1}} \underbrace{\prod_{\forall i: F_i=1} \left(\frac{P(F_i=1|R=1)}{P(E_i=1)P(F_i=1)} P(E_i=1|d) \right)}_{\text{part 2}} \quad (12)$$

where \mathbf{E}' is the set of all possible vectors with $E_i = 1 \forall F_i = 1$. Part 2 of the Eq. (12) is same for all the vectors in \mathbf{E}' . So, we can rewrite Eq. (12) as

$$P(R = 1|d, q) \propto_R \prod_{\forall i: F_i=1} \left(\frac{P(F_i=1|R=1)P(E_i=1|d)}{P(E_i=1)P(F_i=1)} \right) \left(\sum_{\mathbf{E}'} \prod_{\forall i: F_i=0} \left(\frac{P(E_i|F_i=0, R=1)P(F_i=0|R=1)}{P(E_i)P(F_i=0)} P(E_i|d) \right) \right) \quad (13)$$

Eq. (13) is a relevance ranking function under Logical Inclusion with *Query eliteness assumption*.

We can further simplify the Eq. (13) by ignoring the terms that are rank independent and the terms with non-elite properties of the query in Eq. (13). This is similar to ignoring the terms that are not present in query in [8, 5, 3]. Applying a logarithm transform results in the following ranking function:

$$P(R = 1|d, q) \propto \sum_{\forall i: F_i=1} \log \frac{P(E_i=1|d)}{P(E_i=1)} \quad (14)$$

Relationship with IDF: One of the interesting by-products is that the above formula provides a yet another theoretical justification of IDF (inverse document frequency) as scoring function [?]. To see this, let us assume that the eliteness of an term-description elite property is elite to a document if the term is present in the document and non-elite otherwise. Then, the probability of the eliteness of elite property in the collection is to be $P(E_i = 1) := \frac{n_{ci}}{N}$, where n_{ci} is the number of documents in the collection with the term associated with the i^{th} term-description property and N is the total number of documents in the collection. From the above assumption, $\mu_i(0) = 0$; $P(E_i|d) := 1$ if $F_i = 1$ (term i is in the query). By substituting them in the ranking function Eq. (14) we get

$$P(R = 1|d, q) \propto \sum_{\forall i: F_i=1} \log \frac{1}{p_i} = \sum_{\forall i: F_i=1} \log \frac{N}{n_{ci}} \quad (15)$$

The ranking function in Eq. (15) is simply a function of IDF values of the query terms. Basically it says that the IDF score function relies on the assumption that a term is elite if it occurs in the document. This is different from the explanation provided by the Robertson-Spärck Jones model, where an explicit assumption that the whole collection is non-relevant set is needed [8].

4 A Simple Ranking Function and its Experiments

In order to test the theory, we have implemented the basic ranking function in Eq. (14). To compute Eq.(14), we first define a single elite property for each term in the vocabulary and use this set of properties as a complete set of elite properties.

4.1 Estimating Eliteness Probabilities

From the Eliteness Hypothesis, we know that, the occurrence of a term in a document has a stochastic element associated with the eliteness of its corresponding term-description property. Therefore, we compute the probability of i^{th} term-description property being “elite” for the given document d as

$$P(E_i = 1|d) \equiv P(E_i = 1|tf_i) \quad (16)$$

where tf_i denotes the term frequency associated with i^{th} term-description property in document d . From the hypothesis, some documents in the collection are “elite” to an elite property and others are non-elite. And, tf_i follows one distribution in elite set of documents and another distribution in non-elite set of documents. Therefore, we can draw a probabilistic inference about the eliteness of a term-description property from its associated term’s term frequency in the document.

By applying Bayes’ rule to $P(E_i = 1|tf_i)$, we get

$$P(E_i = 1|tf_i) = \frac{P(tf_i|E_i = 1)P(E_i = 1)}{P(tf_i)} = \frac{P(tf_i|E_i = 1)P(E_i = 1)}{\sum_{E_i \in \{0,1\}} P(tf_i|E_i)P(E_i)}. \quad (17)$$

By using Eq. (17) we compute ranking function value for each property in Eq. (14) as follows

$$\frac{P(E_i = 1|tf_i)}{P(E_i = 1)} = \frac{P(tf_i|E_i = 1)}{\sum_{E_i \in \{0,1\}} P(tf_i|E_i)P(E_i)}. \quad (18)$$

For simplicity, we consider only term-description properties, and use query terms to represent elite properties of the query ($F_i = 1$ when $q_i = 1$). By substituting Eq. (18) back in Eq. (14), we get the following ranking function

$$\begin{aligned} P(R = 1|d, q) \\ \propto \sum_{\forall i: F_i = 1} \log \frac{P(E_i = 1|d)}{P(E_i = 1)} &= \sum_{\forall i: F_i = 1} \log \frac{P(tf_i|E_i = 1)}{\sum_{E_i \in \{0,1\}} P(tf_i|E_i)P(E_i)} \end{aligned} \quad (19)$$

To compute Eq. (19), we adopt the 2-Poisson model in [2, 10].

4.2 The 2-Poisson mixture and the EM algorithm

According to the 2-Poisson model hypothesis [2, 10], if the elite property is a term-description property then its associated term's frequency follows a Poisson distribution in the elite set of documents

$$P(tf_i|E_i = 1) := e^{-\mu_i(1)}\mu_i(1)^{tf_i} \quad (20)$$

and another Poisson distribution in the non-elite set

$$P(tf_i|E_i = 0) := e^{-\mu_i(0)}\mu_i(0)^{tf_i} \quad (21)$$

where $\mu_i(1)$ and $\mu_i(0)$ are the two Poisson means. The mixing probability $P(E_i = 1) := p_i$ is an additional parameter. The 2-Poisson mixture can be written as

$$\begin{aligned} P(tf_i|\theta) &= \sum_{E_i \in \{0,1\}} P(tf_i|E_i)P(E_i) \\ &= p_i \frac{e^{-\mu_i(1)}\mu_i(1)^{tf_i}}{tf_i!} + (1 - p_i) \frac{e^{-\mu_i(0)}\mu_i(0)^{tf_i}}{tf_i!} \end{aligned} \quad (22)$$

where θ is unknown parameters of the mixture and $\theta = (\mu_i(1), \mu_i(0), p_i)$. Different with [2, 10], we obtained the optimal maximum likelihood parameters using the Expectation Maximization algorithm [1], where the likelihood function is

$$\begin{aligned} \Lambda(tf_i; \theta) &= \prod_{n=1}^N P(tf_i^n|\theta) \\ &= \prod_{n=1}^N \left(p_i \frac{e^{-\mu_i(1)}\mu_i(1)^{tf_i^n}}{tf_i^n!} + (1 - p_i) \frac{e^{-\mu_i(0)}\mu_i(0)^{tf_i^n}}{tf_i^n!} \right) \end{aligned} \quad (23)$$

where N is the total number of documents and tf_i^n denotes the term frequency of the term associated with i^{th} term-description elite property in n^{th} document (d_n) of the collection.

The parametric estimation problem can be defined as

$$\hat{\theta} = \arg \max_{\theta} \Lambda(tf_i; \theta) \quad (24)$$

where the parameters can be obtained by maximizing the log-likelihood function Λ . More specially, we compute the partial derivatives of the logarithm of the likelihood function with respect to $\mu_i(1)$, $\mu_i(0)$ and p_i , and setting them to zero. The logarithm of the likelihood function $\Lambda(tf_i; \theta)$ is

$$\lambda(tf_i; \theta) := \sum_{n=1}^N \log \left(p_i \frac{e^{-\mu_i(1)}\mu_i(1)^{tf_i^n}}{tf_i^n!} + (1 - p_i) \frac{e^{-\mu_i(0)}\mu_i(0)^{tf_i^n}}{tf_i^n!} \right) \quad (25)$$

By setting

$$\frac{\partial \lambda}{\partial \mu_i(1)} = 0, \quad \frac{\partial \lambda}{\partial \mu_i(0)} = 0 \quad (26)$$

we get,

$$\mu_i(1) = \frac{\sum_{n=1}^N P(E_i = 1|tf_i^n) tf_i^n}{\sum_{n=1}^N P(E_i = 1|tf_i^n)} \quad (27)$$

and

$$\mu_i(0) = \frac{\sum_{n=1}^N P(E_i = 0|tf_i^n) tf_i^n}{\sum_{n=1}^N P(E_i = 0|tf_i^n)}. \quad (28)$$

By setting

$$\frac{\partial \lambda}{\partial p_i} = 0 \quad (29)$$

we get,

$$p_i = \frac{1}{N} \sum_{n=1}^N P(E_i = 1|tf_i^n) \quad (30)$$

where $P(E_i = 1|tf_i^n)$ can be computed using Eq. (17).

In order to find the optimal parameters, we start with an initial estimates of $p_i^{(0)}$, $\mu_i(1)^{(0)}$, and $\mu_i(0)^{(0)}$, and then the better estimates can be found by iterating through the following E- and M-step [1].

1. E-step:

$$P^{(k+1)}(E_i = 1|tf_i) = \frac{P^{(k)}(tf_i|E_i = 1)p_i^{(k)}}{\sum_{E_i \in \{0,1\}} P^{(k)}(tf_i|E_i)P(E_i)}. \quad (31)$$

2. M-step:

$$\mu_i(0)^{(k+1)} = \frac{\sum_{n=1}^N P^{(k)}(E_i = 0|tf_i^n)tf_i^n}{\sum_{n=1}^N P^{(k)}(E_i = 0|tf_i^n)}. \quad (32)$$

$$\mu_i(1)^{(k+1)} = \frac{\sum_{n=1}^N P^{(k)}(E_i = 1|tf_i^n)tf_i^n}{\sum_{n=1}^N P^{(k)}(E_i = 1|tf_i^n)}. \quad (33)$$

$$p_i^{(k+1)} = \frac{1}{N} \sum_{n=1}^N P^{(k)}(E_i = 1|tf_i^n). \quad (34)$$

where k is the step number. We measure the difference of the perplexity between two consecutive models of the probability distribution of $P(E_i = 1|tf_i)$ to decide when shall we stop the training. By substituting the estimated parameter values in Eq. (14) we reach the final ranking function

$$P(R = 1|d, q) \propto \sum_{\forall i: F_i=1} \log \left(\frac{e^{-\mu_i(1)} \mu_i(1)^{tf_i}}{p_i e^{-\mu_i(1)} \mu_i(1)^{tf_i} + (1 - p_i) e^{-\mu_i(0)} \mu_i(0)^{tf_i}} \right) \quad (35)$$

Model	MAP	MRR	Recall@1000
BM25	0.250	0.638	0.6634
Language Model with JM smoothing	0.238	0.4816	0.658
Language Model with Dirichlet prior	0.2539	0.6376	0.6694
Unified Model	0.2553 (0.2266*)	0.607 (0.6513*)	0.6659

Table 1: Performance on the TREC-8 ad hoc task data collection.

4.3 Document Length Normalization

One of the issues in the 2-Poisson model is that it assumes a fixed document length for all the documents [2, 11]. Generally, it is not a valid assumption. Two hypothesis were proposed to explain the varied document lengths in the collection namely Scope Hypothesis and Verbosity hypothesis [11, 12]. Similarly, in this paper we assume the verbosity hypotheses, i.e. lengthy documents covers the same scope of short document but uses more words, and incorporate it into our ranking function by modifying the term frequency as

$$tf = tf \left(b + (1 - b) \frac{avgDL}{DL} \right), \quad (36)$$

where $avgDL$ is average document length in the collection, DL is the document length and $b \in [0, 1]$. The above formula penalizes the term frequency in lengthy documents based on the b value.

4.4 Experiments

For the experiment, we employed the TREC-8 ad hoc task collection and topics to evaluate the ranking function in Eq. (35) against Mean Average Precision (MAP), Mean Reciprocal Rank (MRR) and Recall@1000 metrics. We initialized the parameters with the data collection statistics. The initial value of p_i was set to be the portion of documents collection with the term (same as the assumption used to show the ranking function relationship with IDF), and $\mu_i(1)$ was initialized with the average number of times the term appeared in document with its term frequency more than one. We used a minuscule value to initialize $\mu_i(0)$. The assumption is that the average term frequency of a term associated with the term-description elite property in a document approaches zero if it is “non-elite” to the document.

We have run tests and compared the performance of our method with BM25 [9] and Language Models [5] (with the Jensen-Mercer smoothing and Dirichlet prior). Table 1 shows that our method outperforms the baselines when directly tuning the parameters with respect to the metrics. The improvement over MRR is significant. We also found that MAP and MRR cannot be optimized simultaneously and this confirmed the theoretical argument about the trade-off between MAP and MRR in [13].

As discussed, the parameter values were learned by boosting the initial value of $\mu_i(1)$ (i.e. the average term frequency in elite set of documents.) and by multiplying with an integer n , i.e. $\mu_i(1)^{(0)} = n \cdot \mu_i(1)^{(0)}$. To see the effect of the document length normalization as well as the initialization of $\mu_i(1)$, we fixing the




Fig. 1: Different metrics performances for different values of b for a fixed value of $n = 3$.

n , and vary b and vice-versa. The graph in Fig. 1 shows that the document length normalization indeed has an effect on performance and the best performance was achieved at $b = 0.64$ for $n = 3$.

5 Conclusion and Future work

We have proposed a theoretical unified retrieval framework based on the Eliteness Hypothesis and derived a unified probabilistic relevance ranking function. Our initial experiments with a basic ranking function demonstrated that the approach is indeed encouraging. One of the problems with our current estimation is that the EM algorithm does not always converge to the global maximum. In the future, we would like to test the model by estimating the eliteness probabilities with Bayesian methods and also remove Query eliteness assumption to incorporate the uncertainty in the eliteness of query elite properties and test the model.

References

1. A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *J. Royal Statistical Society, Series B*, 1977.
2. S. Harter. A probabilistic approach to automatic keyword indexing. *Journal of the American Society for Information Retrieval Science*, 1975.
3. J. D. Lafferty and C. Zhai. Document language models, query models, and risk minimization for information retrieval. In *SIGIR*, pages 111–119, 2001.
4. M. E. Maron and J. L. Kuhns. On relevance, probabilistic indexing and information retrieval. *J. ACM*, 1960.
5. J. M. Ponte and W. B. Croft. A language modeling approach to information retrieval. In *SIGIR*, 1998.
6. S. Robertson. The unified model revisited. In *Presented at SIGIR 2003 Workshop on Mathematical/Formal Models in Information Retrieval*, 2003.
7. S. Robertson, M. E. Maron, and W. S. Cooper. Probability of relevance: a unification of two competing models for document retrieval. *Information Technology: Research and Development*, 1982.
8. S. Robertson and K. Spark Jones. Relevance weighting of search terms. *Journal of the American Society for Information Science*, 1976.

9. S. Robertson and H. Zaragoza. The probabilistic relevance framework: BM25 and beyond. *Found. Trends Inf. Retr.*, 2009.
10. S. E. Robertson, C. J. van Rijsbergen, and M. F. Porter. Probabilistic models of indexing and searching. In *SIGIR*, 1980.
11. S. E. Robertson and S. Walker. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR*, 1994.
12. A. K. Singhal. *Term weighting revisited*. PhD thesis, Ithaca, NY, USA, 1997. UMI Order No. GAX97-14899.
13. J. Wang and J. Zhu. On statistical analysis and optimization of information retrieval effectiveness metrics. In *SIGIR*, 2010.